OpenGIS Catalog Service Interfaces RFI Technical Response

National Imagery and Mapping Agency

December 16, 1996

OpenGIS Project Document 96-017

Table of Contents

1. Preface

1.1 Submitting Organization

We are pleased to submit this technical response to the OGC Request 3, OGIS Services Working Group, A Request for Information: OpenGIS Catalog Service Interfaces (OpenGIS Project Document Number 96-017):

1.2 Submission Contact Points

On behalf of the National Imagery and Mapping Agency, all questions about the submission should be directed to:

Mr. Ronald D. Carlson
The MITRE Corporation
1820 Dolley Madison Boulevard
Mail Stop W967
McLean VA 22102
phone: +1 703-883-6971
email: ronc@mitre.org

2. Overview

On 1 October 1996 the Department of Defense (DoD) formally established the National Imagery and Mapping Agency (NIMA). NIMA's mission is to provide timely, relevant, and accurate imagery, imagery intelligence, and geospatial information in support of national security objectives. NIMA's United States Imagery and Geospatial Information System (USIGS) is the umbrella term for the suites of systems that had formerly been called the Global Geospatial Information and Services (GGIS), and the United States Imagery System (USIS). A very important part of the USIGS are the system components, network resources, and interoperability mechanisms to provide customers, whether internal or external to NIMA, with the catalog services necessary to search, access, and retrieve available holdings of geospatial and imagery data. NIMA's current plans are to accomplish this through the implementation of a set of library servers that will meet these goals.

Part of the NIMA vision for USIGS involves the establishment of a database or suite of databases that will provide the user with the right information at the right time with the needed accuracy, currency, and completeness. According to "Global Geospatial Information and Services (GGIS) for the Warrior" [1], NIMA will establish a gateway making its digital data holdings available to a broadening base of customers by taking advantage of the decreasing costs and increased capabilities of modern computing and telecommunications resources and equipment. These data holdings will be accessible to users through an interface supporting the necessary catalog services to access NIMA product inventories and metadata.

To achieve the goals and objectives embodied in the USIGS vision, NIMA has committed to using a technical reference model as framework for the development and implementation of the USIGS. A technical reference model is a high-level representation of a system's architecture that enables people to agree on definitions, build a common base of understanding, and identify and resolve crucial issues. A technical reference model is not a specific system design, but rather, establishes a standard vocabulary for the conceptual modules, services, and interfaces that comprise an architecture, in order to provide a context for analyzing portability, scalability, interoperability, and other technological issues.

NIMA has selected the Common Imagery Interoperability (CIIF) Reference Model [2] as the reference model to be used for USIGS. The CIIF Reference Model has recently been revised to encompass additional geospatial information constructs beyond its previous emphasis on imagery. The USIGS architecture in general, and CIIF API standards in particular, relate to emerging architectural concepts and standards connected with the Defense Information Infrastructure Common Operating Environment (DII/COE). The CIIF Reference Model shown in Figure 2-1 was derived by combining selected features of the current DII/COE reference model, the Object Management Group's Object Management Architecture (OMA) reference model, and the Intelligence Community's draft reference model.

This composite model lays the groundwork for migrating to a fully object-oriented distributing computing architecture, while recognizing that object-oriented and non-object based implementations will need to coexist for some time to come. The key concepts contained in this model include:

- Use of Standard Application Program Interfaces to provide a fundamental enabling mechanism for better integrating the diverse collection of client applications, server applications, and operating system services that comprise the USIGS environment.
- Promulgation of an open architecture that will use these standard APIs to facilitate the exchange of information and sharing of services among all manner of Mission-Specific Applications.
- Establishment of a fundamental collection of Infrastructure Services (accessible via a standard API), which provide the basic capabilities needed to implement a distributed computing system architecture
- Definition of a collection of Common Support Applications and Common Facilities, which provide a variety of ancillary services and utilities that are useful (and at time quite necessary) for building and standardizing distributed systems.

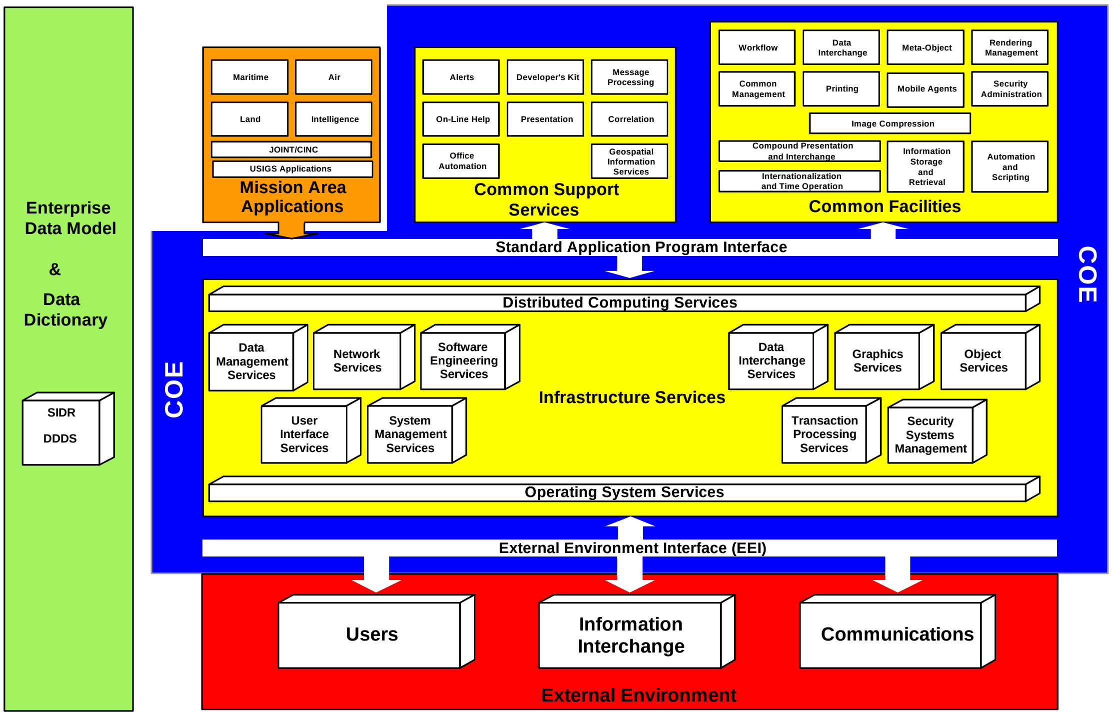

Figure 2-1: CIIF Reference Model

The architecture of the USIGS Software Reference Model, shown in Figure 2-2, is based on the Long-Term CIIF Reference Model which is adapted from the Object Management Group's (OMG) Object Management Architecture (OMA). Two of the components included in the CIIF Reference Model are particularly important in regard to understanding its connection with the Long-Term CIIF Reference Model, as well as between this model and the details of the OMG OMA architecture. First, the Mapping, Charting, Geodesy and Imagery (MCG&I) capabilities of the USIGS are contained in the Geospatial Information Services component of the Common Support Applications in the CIIF Reference Model, which is called the GIS Domain Common Services in the USIGS Software Model. It is here that the domain interface standards that are the primary focus of the CIIF may be found. Second, the Object Services component of the Infrastructure Services contains the Object Service Interfaces that are fundamental to the OMG OMA architecture.

The USIGS Software Reference Model classifies the components, interfaces, and protocols that comprise an object system. This long-term model represents the long-term goal of a fully-realized object oriented environment. It is envisioned that as the technical architecture migrates towards purely object oriented designs, the systems architectures will evolve towards compliance with this model. The Infrastructure Services identified in the CIIF Reference Model (see Figure 2-1) are anticipated to be replaced by Object Service Interfaces and Common Facilities. The USIGS Software Reference Model has six key components:

- **Distributed Computing Services** - enables software objects to make and receive requests and responses within a distributed environment.
- **Object Services** - a collection of fundamental services (interfaces and objects) that provide basic functions for using and implementing other software objects.
- **Common Facilities** - a collection of higher-level services that are broadly usable by many applications.
- **GIS Domain Common Services** - standard interfaces that promote object-based interoperability within the imagery and mapping community or application domains.

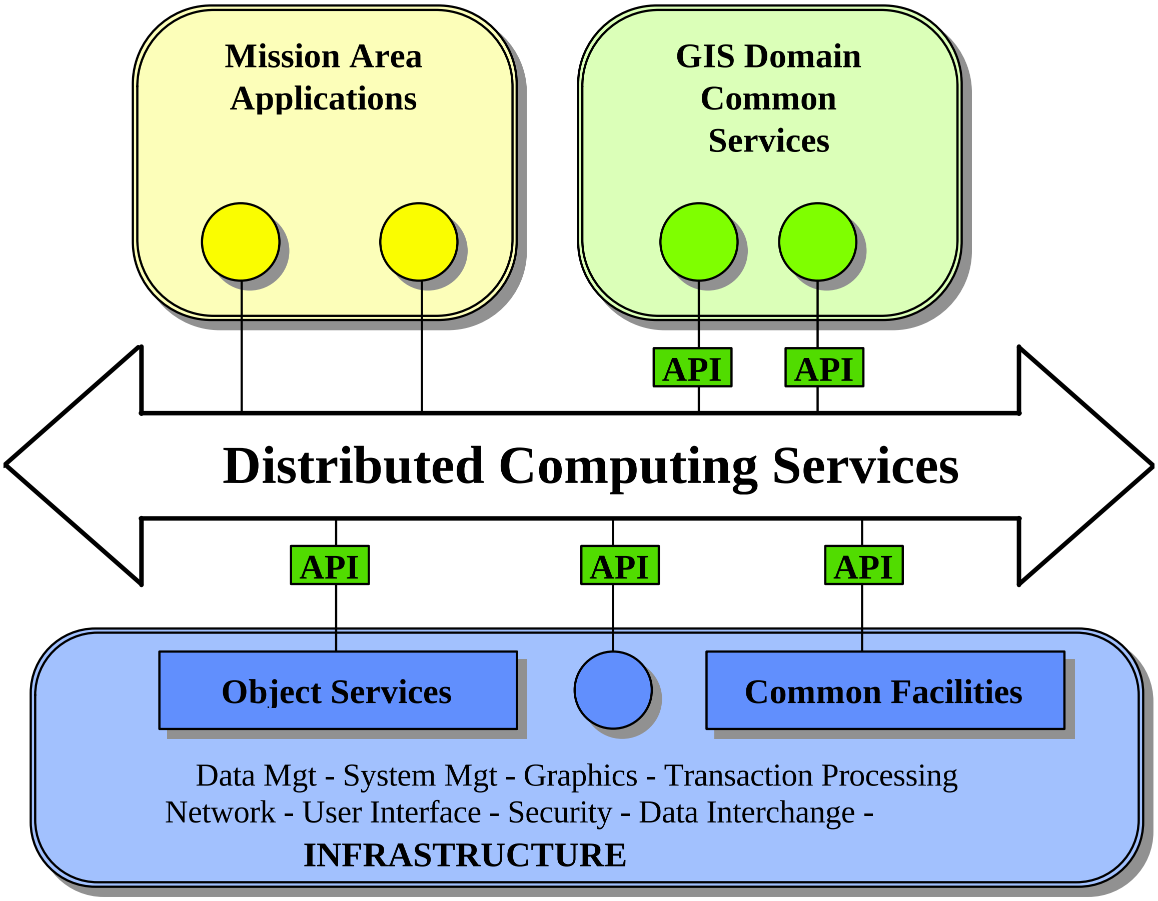

Figure 2-2: USIGS Software Reference Model

- **Mission Area Applications** - software objects specific to the USIGS, including particular commercial products or end-user systems.
- **Interface Definition Language (ISO IDL)** - A formal language (Information Technology -- Open Distributed Processing -- Interface Definition Language, ISO Draft Standard 14750 [3]) used to define the interfaces between interoperable software objects. The benefits of ISO IDL are discussed in Section 3.12.

At a more detailed level, Common Facilities and Geospatial Information Services define those interfaces and uniform sequencing semantics that are shared across applications in such a way as to make object-oriented distributed computing applications much easier to create. Common Facilities and Geospatial Information Services comprise both generic facilities and domain-specific specifications. Examples of the kinds of inter-application services provided by Common Facilities and Geospatial Information Services include object cataloging and browsing, help facilities, object rendering, printing and spooling, and objects which implement generic business rules for the imagery and mapping industry.

The roles, uses, and definitions of individual Common Facilities and Geospatial Information Services have the potential to evolve over time. Geospatial Information Services that are used extensively in a variety of similar applications may gradually change from being domain-specific to being generic in character. Moreover, services that are offered across multiple application domains are good candidates for incorporation into future versions of the Common Facilities. Discovery of such commonalities will be a favorable indicator of the maturation of these standards.

The boundaries separating Common Facilities and Geospatial Information Services from Mission-Specific Applications (in the one direction) and from Object Service Interfaces (in the other direction) are therefore not fixed and immutable, but rather, are a reflection of the state-of-the-art in object system technology. As experience in a particular application domain advances, areas of potential new Geospatial Information Services or Common Facilities will be discovered and defined–just as evolving system infrastructures will gradually incorporate pieces of the Common Facilities into their basic Object Service Interface offerings.

Operations provided by the Object Service Interfaces component of the CIIF Reference Model are expected to serve as key building blocks for Common Facilities, Geospatial Information Services, and Mission-Specific Applications. Common Facilities and Geospatial Information Services, in turn, provide higher-level interoperable interfaces that can be specialized for particular Mission-Specific Applications. The practical application of these various levels of standard interfaces and services makes extensive use of (and in fact, depends upon) the object-oriented concept of inheritance. Compiler-based support for these inheritance processes has been built into ISO's Interface Definition Language (IDL), via its various standard language mappings. Inheritance facilitates the standardization of interfaces, promotes interoperability between objects conforming to the base standard, and enables the design of consistent interfaces between otherwise disparate object types.

In non-object-oriented software architectures, a system's Application Program Interface (API) is typically defined by a fixed, monolithic, interface structure. The IDL-based API that is intrinsic to the CIIF Reference Model, on the other hand, is fully modular–the developer of an application object can pick and choose exactly those interface elements that are needed to fulfill design objectives. In other words, in contrast to the case for the traditional style of API, the object-oriented APIs provided by IDL-specified Geospatial Information Services and Common Facilities are extensible, customizable, and subset-able. This high degree of adaptability will facilitate the controlled evolution of such object-based standards, thereby helping to ensure the long-term viability of the technology. These concepts are illustrated in the software component reuse diagram in Figure 2-3 below:

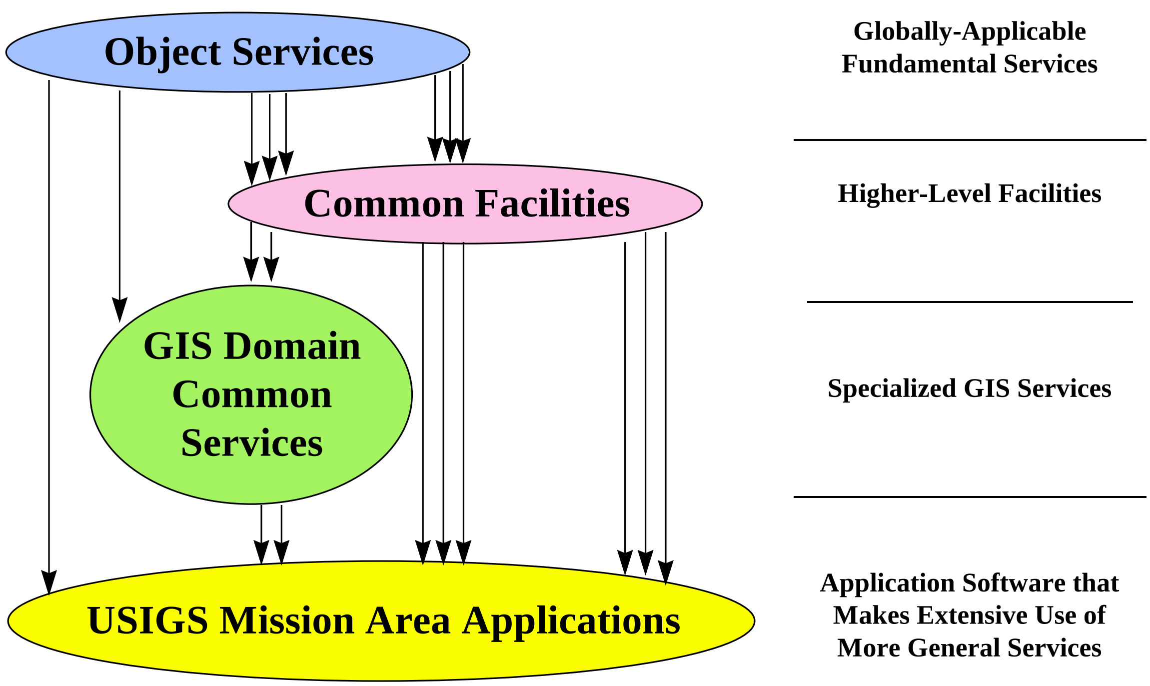

Figure 2-3 Software Component Reuse

The basic philosophy on which the CIIF Reference Model [2] is based embraces the object-oriented paradigm of software development. The CIIF architecture is based on a network-centric concept of systems interoperability. This concept reflects the growing importance of widely distributed computer interactions for running the modern business or governmental enterprises. It also reflects the growing trend for software vendors to include support for advanced telecommunications capabilities in their operating system and application products. The design of a CIIF-compliant facility should naturally accommodate extensions via iterative development processes. The specification for a facility should make clear how such extensions are to be defined (via inheritance, delegation, etc.). Facilities should use the multiple inheritance capability of ISO IDL wherever appropriate, since multiple inheritance eases the task of defining interfaces that address varied purposes. NIMA fully supports this approach and, with it, the extensibility and specialization benefits inherent in this approach.

In executing its technology development process for the definition of geospatial catalog service interfaces, OGC should strongly consider an approach that is consistent with the framework established by the guiding philosophy and technical reference model concepts summarized in this RFI response and presented in detail in the accompanying reference documents. In adopting such an approach, OGC should focus on facilitating the development of the GIS Domain Common Services portion of the CIIF Reference Model through its unique relationship with the geospatial software application vendor community. A focus on the Mission Area Applications to be created using the GIS Domain Common Services to meet the requirements of the USIGS is the responsibility of NIMA. Its is through the geospatial APIs to be developed through the OGC technical development process that NIMA expects to tie USIGS Mission Area Applications together.

In responding to this RFI, we found most of the request items to be too narrowly focused, though we have responded as appropriately to each request item as possible based on our experience in developing the CIIF Reference Model and its associated specifications. Our response takes the position that what is really being asked for relates to a broad range of digital library services for imagery and geospatial information holdings rather than the more limited topic of catalog services. We found it necessary to alter the internal structure of the RFI response by combining our responses to RI2 and RI3 since we strongly believe that identified sets of service components comprising a catalog and their interdependencies are two topics that, to be discussed appropriately, must be addressed jointly and cannot be separated from one another.

3. Technical Response

3.1 RI1 The Definition Of A Geospatial Catalog Facility

As mentioned in Section 2, we take the position that the entire spectrum of digital library services should be of interest to OGC. To gain insight into a digital catalog facility and the complete range of library services offered, such as ingest, maintenance, management, and dissemination capabilities, it is instructive to examine the work being done by ISO to get a more complete picture of the digital catalog and archival problem.

An ISO Working Group has convened under the leadership of NASA to create a proposed standard for digital archives. The draft document "Reference Model For An Open Archival Information System (OAIS), Version 7.0" [4] defines a minimum set of responsibilities for the recognition of an OAIS archive and it provides a framework for understanding the challenges associated with permanent or indefinite long-term information preservation, especially for digital information related to space missions such as imagery and other remote sensing products. The framework provided by the reference model facilitates the description and comparison of the architectures and operations of existing and future information-preserving archives.

A major purpose of this reference model is to facilitate a much wider understanding of what is required to preserve information permanently or for the indefinite long term. It is expected that this reference model, by establishing minimum requirements and mandatory responsibilities for an OAIS archive along with a set of archival concepts, will provide a common framework from which to view archival challenges, particularly as they relate to digital information. This should enable more organizations to understand the issues and to take proper steps to ensure long term information preservation. It should also provide a basis for more standardization and therefore a larger market that vendors can support in meeting archival requirements.

The OAIS model recognizes the already highly distributed nature of digital information holdings and the need for local implementations of effective policies and procedures supporting information preservation. Though no implementation is specified or mandated, standards developers are expected to use this reference model as a basis for further standardization. A large number of related standards may possibly emerge from this reference model upon its completion, such as minimal attribute set standards, adopted conformance to the IEEE Reference Model for Open Storage Systems Interconnection [5] for some or all APIs, or adoption of selected profiles for access to digital collections under development by groups related to the Z39.50 protocol [6] and [7].

The concept of an Archival Information Package (AIP) is used in the OAIS reference model to represent archival information at the highest level to provide a concise way of referring to a set of information that has, in principle, all the qualities needed for permanent, or indefinite-long term, preservation of a designated information object. Within the AIP, this designated information object is called the Content Information. The Content Information is that information which is the primary object of preservation. Also within the AIP is a set of information called the Preservation Description Information. The Preservation Description Information is that information which is necessary to adequately preserve the particular Content Information with which it is associated so the Content Information remains sufficiently useful to independent parties indefinitely into the future. This information includes lineage (i.e., provenance) information; reference information that identifies and describes the Content Information; context information that relates the Content Information to its environment; authentication mechanisms; and catalog metadata.

The environment model of an OAIS, shown in Figure 3-1, differs somewhat from the model presented in the catalog use case scenario shown in Figure 1 of the Catalog Service Interfaces RFI. Whereas the catalogs and storage collections are shown as separate objects in the RFI, the OAIS treats the archived AIP and its accompanying metadata as a single object when viewed from the external environment. The "users" shown in the RFI scenario are referred to as "consumers" in the OAIS environmental model, while the "librarians" shown in the RFI are not considered an external actor in the OAIS environmental model. The day-to-day management functions performed by librarians are considered to be an integral part of the OAIS. The RFI scenario omits the "producer" and "management" actors present in the OAIS model altogether.

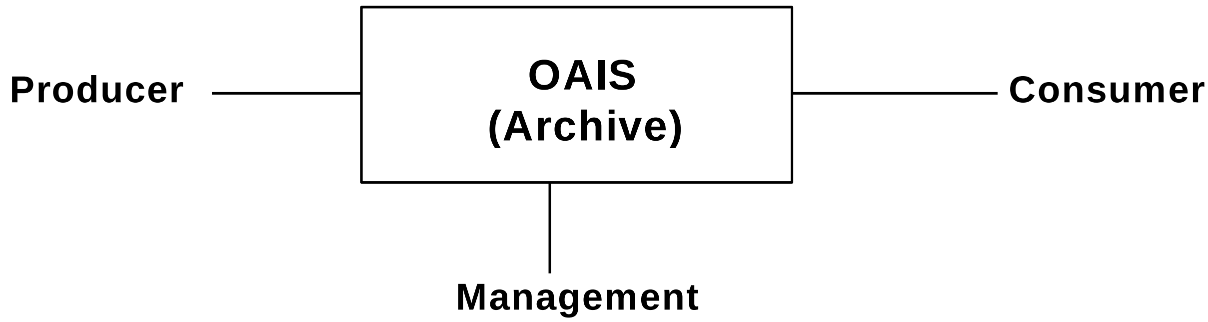

Figure 3-1: The Environmental Model of an OAIS [4]

The OAIS reference model addresses a full range of archival information preservation functions including ingest, storage, access, and dissemination. It also addresses the migration of digital information to new media and forms, the data models used to represent the information, the role of software in information preservation, and the exchange of digital information among archives. It identifies both internal and external interfaces to the archive functions, and it identifies a number of high level services at these interfaces. A first step toward establishing the ability to characterize various types of archives, and their quality, is also addressed.

The OAIS is organized into seven functional areas and related interfaces as shown in Figure 3-2. The lines connecting functions identify communication paths over which information flows in both directions. No connections are shown for Common Services to avoid clutter as this function communicates with all other functions. Similarly, no internal connections are shown to Administration because this function also communicates with all other functions.

Common Services. The Common Services entity provides a single repository for those supporting services in modern, distributed computing applications that are common to one another. Some of these services include interprocess communication, name services, temporary storage allocation, exception handling, security, and directory services.

For large archives, it is particularly important to have a generic policy management service to provide mechanisms for automated policy definition and monitoring for compliance to policy. An important differentiation between archives as defined in this report and other data storage sites is the existence of internal policies and external contracts to ensure the preservation of and access to the archived information over long periods of time. Given the potential size and complexity of digital archives these policies would be impossible to support without automated assistance. Some of the areas where policies are required include costing, media monitoring for degradation, backups, product identification, federation and interoperability with other archives, and preservation of information under impending archive dissolution.

Ingest. The Ingest entity provides the services and functions to accept and validate information objects and prepare them for storage and management within the archive. These services and functions include:

- *Staging*. The staging function makes physical storage space available for a data submission session. The archive staff must allocate the appropriate storage capacity or devices to the data preparer. The data may be electronically transferred via ftp to the staging area; loaded from media submitted to the archive or simply mounted (e.g. CD-ROM) on the archive file system for access. The data submission formats and procedures must be clearly documented in the archive's data submission manual and the deliverables identified by the preparer in a data management plan or archive delivery plan.
- *Review*. The review function provides a validation of the submission package and assures that the information is understandable to the intended consumer community. The review may be carried out by the archive data engineers and may also involve an outside committee (peer review). The review must verify three things, 1) that the data has been physically transported correctly to the archive staging area; 2) that

the quality of the data meets the requirements of the archive or peer review group; 3) that the documentation and

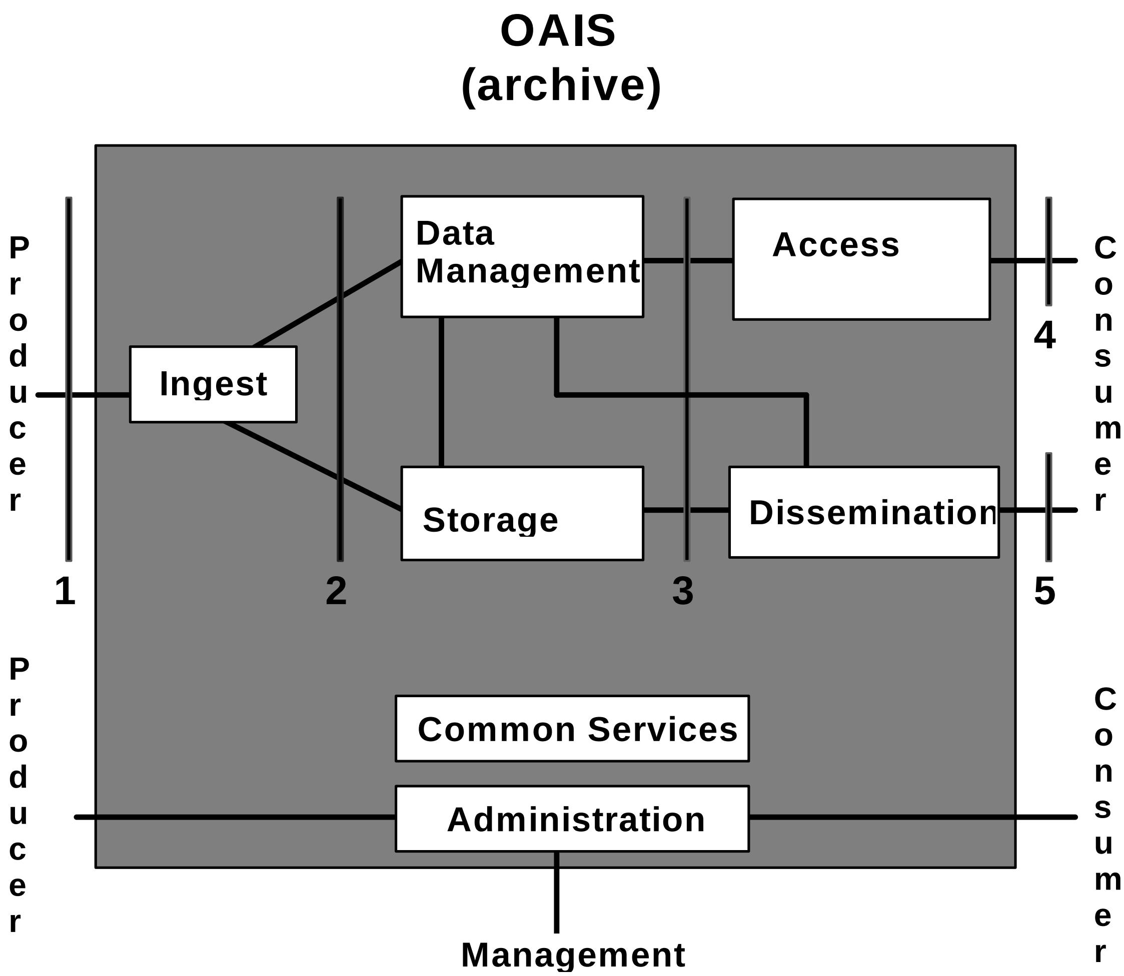

Figure 3-2: OAIS Functional Entities [4]

metadata are sufficient to make the data understandable to the target user community. The formality of the review will vary depending on internal archive policies. The review process may determine that some portions of the submission package are not appropriate for inclusion in the archive and must be redone or excluded.

- *Conversion*. The conversion function transforms the submitted objects into objects and descriptions that conform to the internal data standards of the archive. It should be a high priority of the archive to see that the majority of submission packages follow archive format and content standards as specified in a Data Submission Manual. Most conversions involve format. The archive is assumed to have a pool of standard utilities to support the conversion of submitted data or metadata objects to internal format.
- *Extraction*. The extraction function generates metadata needed to support one or more access views to the archive data. This may include the production of special versions of a data object to support on-line access (thumbnail or browse images, for example), or production of special formats of data products or documentation (e.g. Acrobat PDF files).

- *Transfer*. The transfer function moves the data objects from the staging area to the storage area and provides metadata to the data management function for cataloging. This may be either an electronic or physical transfer. The transfer of metadata to the data management entity should include the set of transactions procedures required to update the archive data base . The transfer of data objects to the storage entity will include a storage request detailing the location, media type and volume of data to be stored.
- *Ingest Reporting*. The ingest reporting function provides interaction between the archive staff and the data preparer. The initial confirmation is an acknowledgment of receipt of a submission package. After the review process is completed any liens are reported to the preparer who will then resubmit or appeal the decision. After review completion a final report on the submission session is prepared for distribution to the administration and to the producer.

Archival Storage. The Archival Storage entity provides the services and functions for the storage and retrieval of data objects, both logical and physical, including non-digital media. These services and functions include:

- *Transfer.* The transfer function receives a transfer request and data objects from staging storage and moves the data to permanent storage within the archive. The transfer request should indicate the anticipated frequency of utilization of the data objects to allow the appropriate storage devices or media to be selected for the data objects. The transfer function will select the media type, prepare the devices or volumes and perform the physical transfer to the storage volumes.
- *Hierarchy management.* The hierarchy management function positions data objects and collections on the appropriate media based on usage statistics or access requirements determined by administration. Hierarchy management procedures must provide policies for the use of on-line versus near-line versus off-line storage. To the extent possible the hierarchy management function should be automated based on usage statistics generated by the storage reporting function.
- *Migration.* The migration function is responsible for maintaining the validity of the data objects across media over time. The migration strategy must take into consideration the expected and actual rates of error encountered in various media types and assure that data objects are converted with minimal data loss. The migration function must find a way to maintain the unique attributes of various media types (e.g. tape block sizes, CD-ROM volume information) when migrating to higher capacity media with different storage architectures.
- *Error checking.* The error checking function provides statistically acceptable assurance that no data objects are corrupted during transfer, migration or backup procedures. This function requires that all hardware and software within the archive provide notification of potential errors and that these errors are routed to standard logs that are checked by the storage staff. A standard mechanism for tracking and verifying the validity of all data objects within the archive should also be used. For example, cyclical redundancy checks (CRCs), should be maintained for every individual data file. The storage facility procedures should provide for random verification of the integrity of data objects using CRCs or some other error checking mechanism.
- *Backup.* The backup function provides an automated mechanism for producing a duplicate copy of the archive contents. The backup media should be capable of being physically removed from the archive for storage at a separate facility. Backup requirements are specified in the archive procedures for various types of data.
- *Duplicate.* The duplicate function provides copies of stored data objects for distribution on physical media to requesters on media types supported by the archive. The duplicate request must provide an itemized list of data objects or physical volumes and identify the output media type. The duplicate function must prepare the output media, perform the transfer from storage to the new media and provide logical and physical labels for distribution with the duplicate copy.

- *Storage Reporting*. The storage reporting function provides regular reports to administration summarizing the inventory of media on-hand, available storage capacity in the various tiers of the storage hierarchy, and operational statistics.

Data Management. The Data Management entity contains the services and functions for populating, maintaining, and querying a wide variety of metadata including product related metadata such as data catalogs, directories, inventories, and processing algorithms; and other metadata including information on customer access and security, archive schedules and procedures, and processing history. In addition, this entity supports data administration and data engineering for the Administration entity. Specific services and functions include:

- Enforcing access controls;
- Receiving requests for customized or periodic reports;
- Generating requested reports;
- Providing transaction control for updates;
- Defining schemas; and
- Providing customized views.

Administration. The Administration entity contains the services and functions needed to control the operation of the other Archive entities on a moment by moment basis. Its services and functions include:

- Establishing acquisition plans between providers and the archive;
- Establishing rules and policies governing all aspects of archive operations;
- Scheduling, monitoring , and controlling of interactions among multiple OAIS entities;
- Establishing policies on migration of digital information representations to new representations;
- Authorizing and authenticating new consumers and producer; and
- Directing the handling of special requests.

Access. The Access entity contains the services and functions which make the archival information and externally-available services visible to consumers, accept orders from consumers and provide customer services. These services and functions include:

- Maintaining and advertising lists of provided services, prices, products, packaging, and options;
- Developing and maintaining the user interface portion of the tools that are available for locating data;
- Authorizing customer access;
- Connecting customers to appropriate finders aids;
- Estimating data volumes and costs of prospective orders;
- Recording data orders (both single orders and recurring orders);
- Passing order information to the Dissemination entity;
- Collecting and recording information and statistics on customer accesses and orders; and
- Providing customer service by addressing the customer's unique needs.

Dissemination. The Dissemination entity contains the services and functions used to request information objects from storage, perform requested transformations, and make the resulting information objects available in the manner requested. These services and functions include:

- receiving data orders;
- retrieving the data and metadata necessary to fulfill an order;

- generating any additional data or metadata needed to fulfill an order;
- formatting data and metadata into customer specified format;
- caching large data orders; rescheduling execution of a standing order;
- receiving notifications that data are ready for delivery;
- determining the medium and method of data delivery;
- recording data onto distribution media if necessary; calculating and recording billing information;
- preparing packing lists, bills of lading, and other shipping records as needed;
- initiating and tracking delivery through to receipt of data by the customer; and
- collecting and recording notices and statistics on delivery.

It is recommended that OGC coordinate more closely with the ISO Working Group that is developing the OAIS reference model to leverage its benefits in aligning the OGIS work with other reference model and standardization activities, and influence the development of the other activities as appropriate.

3.2 RI2 Identified Sets Of Service Components Comprising A Catalog

As described in Section 2, we have combined our response to RI2 and RI3 since we strongly believe that identified sets of service components comprising a catalog and their interdependencies are two topics that, to be discussed appropriately, must be addressed jointly and cannot be separated from one another.

The Imagery Access Services (IAS) Specification [8] addresses the core interoperability requirements of the USIGS for client access to imagery and geospatial information products. The IAS specification defines the interface requirements for three (3) Common Facilities from the CIIF reference model, which was shown in Figure 2-2. These facilities are:

- Image Access Facility (IAF)
- Catalog Access Facility (CAF)
- Profile and Notification Facility (P&NF)

The IAS Specification is fully consistent with the OAIS Reference Model that was described in the response to RI1, and performs many of the basic functions of a digital library and archival service. The facilities specified in the IAS perform some of the Ingest and Access functions and most of the Dissemination functions described in the OAIS, and are defined in terms of the external interfaces necessary to support these functions. The interfaces to these functions are external in that they are primarily oriented toward supporting the consumer of the stored data, rather than the producer or maintainer of the data via internal interfaces. It should not be implied from this that data production or data management are not important interface types, only that these interfaces are internal in nature, not external. Thus, they are not defined in the IAS Specification.

The interfaces of the IAS facilities are mapped to the OAIS functional entities that they support as shown in Table 3-1. The OAIS functional entities were described in the response to RI1, and the operational interface type for each of these are shown across the top of the matrix. The specific interfaces for each of the facilities in the IAS are shown down the left side. An “X” is placed where the row of each IAS interface intersects with the appropriate OAIS functional entity it fully supports. Where the function is only partially supported, the name of the specific module from the IAS is shown.

The **Catalog Access Facility (CAF)** is one of the facilities that comprise the CIIF as specified in the IAS Specification [8]. As a GIS Domain Service, the CAF comprises generic facilities for inter-application services pertaining to object cataloging and browsing. The CAF provides common software interfaces between an imagery and geospatial library catalog and client applications for the purposes of product discovery, product attribute (metadata) retrieval, product indexing, directory maintenance, and uniform access to database resources. The CAF enables imagery and geospatial information searches, imagery and geospatial information browsing, and catalog updates.

As described in the IAS Specification [8], the Catalog Access Facility provides several basic services, called "operations", that are supported through its interfaces. The interfaces are specified in Interface Definition Language (ISO IDL) according to ISO Draft Standard 14750 [3], which is used to define the interfaces between interoperable software objects. The ISO IDL for these operations is available in the IAS Specification. The user operations supported by the CAF to query and browse available information are listed and described below:

		OAIS FUNCTIONAL ENTITIES						
IAS INTERFACE TYPES	Operational Interface Type	INGEST	DATA MGMT	STORAGE	COMMON SERVICES	ADMIN	ACCESS	DISSEMINATION
IMAGE ACCESS	Server Interface		n/a	n/a	n/a			X
	Image Access Interface	create ( )	n/a	n/a	n/a			X
	Image Product Interface		n/a	n/a	n/a			X
	Image Array Interface		n/a	n/a	n/a			X
	Parameters Interface		n/a	n/a	n/a			X
PROFILE & NOTIFICATION	Profile & Notification Interface		n/a	n/a	n/a		X	X
CATALOG ACCESS	Catalog Access Interface		n/a	n/a	n/a		X	

Table 3-1: Mapping of IAS Standard Interfaces To The OAIS Reference Model

- **Boolean Query Operation** - enables ordinary catalog search queries by accepting Boolean Query Syntax expressions as input and returning a set of query hits matching the expression.
- **Polygonal Query Operation** - enables ordinary catalog search queries by supplementing Boolean queries with the specification of a polygonal shape. Image products that overlap any portion of the polygon will satisfy the query if the product attributes also satisfy the Boolean query expression.
- **Elliptical Query Operation** - enables Boolean catalog search queries combined with the specification of an elliptical shape. The ellipse is defined by its center point, major and minor axes, and azimuth of rotation from North of the major axis, and it returns image products which provide coverage of any portion of the ellipse while satisfying the Boolean Query Syntax expression.
- **Point Query Operation** - supplements Boolean queries by specifying a geographic point. Image products returned by this query combine the Boolean Query Syntax expression and coverage of the specified point.
- **Get Results Operation** - accesses invocations of Catalog Access Facility operations unable to return entire sets of query hits in the allocated area for query results. A queried value is returned and used with the "get_more_results" operation to obtain the remaining query hits.

- **Free Results Operation -** notifies the catalog server that the client does not intend to retrieve additional results for the indicated QueryID. This enables the catalog server to free any resources allocated to the indicated QueryID, including any remaining results.

The CAF uses a mechanism called Boolean Query Syntax by specifying conventions on arguments and operation signatures for the purpose of enabling dynamic queries. Boolean Query Syntax is a notation for expressing queries based directly on pre-defined attribute lists. The query is based on an attribute list instead of a particular schema. This approach simplifies the complexity of the client for query generation and avoids constraining the design of the schema or schema view in the server implementation of the CAF. Thus client software is simplified and decoupled from changing community specified attributes for the purposes of querying image catalogs.

The syntax of the Boolean query language is constrained to provide simplicity of query generation and translation without loss of useful capabilities. The net result of a simplified client approach is that the server that must handle additional processing requirements that add to the server's complexity. For instance, the server must provide the processing capabilities that are required to translate queries into the appropriate schema and query language syntax.

This approach is based on the attribute/value pair user interface paradigm that drives the generation of Boolean queries. In this paradigm Boolean query syntax may be constrained to a Boolean product-of-sums, where most of the sums are expressed as simple attribute relations. When this is not the case, a logical sum expression is the result based on the same attribute. This is specified in Backus-Naur Form in the IAS Specification, following the same conventions as used in the ISO IDL technical reference. These rules are augmented by the following constraints:

- Constant expressions include the options defined in SQL92, except as otherwise noted.
- The <Attribute> contained in each relation with a sum expression are the same attribute. The operators are limited to "=", "<>", "like", and "not like" within sum expressions.
- Wildcard expressions are allowed using the character "%" to denote a match with 0 or more characters.
- "Like" and "not like" are the only operators that are used for text expressions and the only operators that support wildcards.

All allowable queries are generated from a Boolean Query. A Product Expression is a logical sum of products, i.e. a series of expressions that are ANDed together. A Sum Expression is either a simple attribute relation or a series of simple relations that are ORed together. A Relation Expression is a simple relationship between a particular attribute and a constant value. Constant values may be integers, floating point, strings, as well as additional options permitted in the SQL92 standard.

Interoperability will be limited if clients generate queries that require excessive processing. Therefore each implementation may identify specific attributes that are able to be queried. Additionally, the implementation may also define the relation operators and wildcards that are allowed, as well as other characteristics. These attributes are the metadata that are available as a set of parameters, retrievable through the Parameters interface. These parameters are defined and retrievable from all archived objects. The QueryableAttributes parameter is a list of attributes and the query capabilities that they support. Each attribute has a separate entry in the name-value list, where the attribute name is the "name" field. The attribute's query capabilities are listed in a string-valued list of allowable operators with space separators and the wildcard symbol, if supported.

The CIIF supports the concept of specialization of the catalog service facilities it provides through the **Profile and Notification Facility (P&NF)**, as described in the IAS Specification [8]. The P&NF is a specialization of the CAF for the purpose of posting standing queries to the network for discovery of imagery and geospatial information resources. The P&NF enables clients to create and manage profiles of search parameters for a given community or set of users that are unique to their special area of interest. These profiles serve as standing catalog search specifications that permit users to register their notification preferences so that the facility implementation can detect when new catalog entries are posted or become available that satisfy the user's profile specification.

The IAS Specification [8] uses ISO IDL to describe two interfaces for the P&NF. The first of these interfaces is called the Profile & Notification Interface. This interface reuses the operations and methods, which were previously described, that it inherits from the CAF interface with relatively no change of semantics. The four query methods (boolean_query, polygonal_query, elliptical_query and point_query) are used to submit standing queries and the get_results method allows result sets from standing queries to be retrieved. Return values are retrieved using the get_results operation. User exceptions that may be generated by the queries are returned from the operation invocation that posts the query. In other words, the query arguments are checked before they are accepted for posting. The free_results method differs slightly from its usage in the CAF, where queries and their associated results are considered transitory and freeing a transitory result deletes all results. The P&NF, in contrast, considers queries and results to be persistent. Freeing the results from a standing query only deletes those results that have already been retrieved by the client. This also means that in the P&NF, a QueryId returned by submitting a standing query is valid for an indefinite time. This is different from a QueryId returned from a CAF query, which is only valid during the session in which it was issued by the server.

Beyond those operations used in the CAF, the Profile & Notification Interface defines three (3) additional operations which are listed and described below:

- **List Queries Operation -** The list_queries operation returns a status list of all of the standing queries. The output argument is a QueryStatusList which contains one element for each standing query from this client. The query_id member uniquely identifies each query with respect to the requesting client. The new_results member indicates if there are new query results to be retrieved. That is, new_results equal to TRUE indicates that a query has results that have not been retrieved. For example, if a standing query has 100 results and the client retrieves 10 (using get_results) the new_results flag should still be set to TRUE, because 90 hits remain. This also means that repeatedly checking status (using list_queries) does not change that new_results status. The client is responsible for maintaining the correlation between the QueryId and the details of the query (i.e. query parameters, human-readable description of the query etc.). The standard exception BAD_INV_ORDER (routine invocations out of order) will be returned if the server has not been successfully opened prior to this method being used.

- **Remove Query Operation -** The remove_query operation allows the cancellation of a standing query. The BadQueryId exception is returned from operations if there is no standing query with the indicated query_id. The standard exception BAD_INV_ORDER (routine invocations out of order) will be returned if the server has not been successfully opened prior to this method being used.

- **Request Notification Method** - The request_notification method indicates that the client wishes to be automatically notified of new "hits" against a standing query. The client supplies a QueryId to indicate the standing query of interest. It also supplies an object reference for a PNF_Callback object for the server to notify. On receiving hits against this standing query, the PNF server will invoke the notify method on this object reference. (see PNF_Callback below). The client can also provide a string containing an email address. On receiving hits against a standing query, an email message describing those hits will be sent to that address. Either the object reference or the address parameter may be NULL, indicating that client does not wish to be notified by that mode, but both cannot be NULL, The standard exception BAD_PARAM would be raised if both are NULL. The BadEmailAddress exception would be raised if an invalid email address is supplied. The exception BadQueryId is returned if the QueryId is invalid. The standard exception INV_OBJREF (invalid object reference) is returned if the client submits an invalid object reference for the PNF_Callback.

- **Request Push Operation -** The request_push operation indicates that the client wishes to be sent all images resulting from hits on the standing query identified by query_identifier. The client supplies a QueryId to indicate the standing query of interest . It also supplies an object reference for a PNF_Callback object for the server to contact prior to sending the imagery or geospatial information. When the a PNF server has imagery or geospatial information to be pushed to the client it will invoke the push method on this object reference. (see PNF_Callback below). The exception BadQueryId is returned if the QueryId is invalid. The standard exception INV_OBJREF (invalid object reference) is returned if the client submits an invalid object reference for the PNF_Callback.

The second interface in the P&NF provides a set of callback interfaces and is referred to as the PNF_Callback Interface. The PNF_Callback interface is implemented by any client that wishes to be automatically notified of standing query hits or to have images automatically pushed to it**.** Clients that do not implement this interface cannot use the request_push method or the callback mode of the request_notification method. They may use the email mode of request_notification. The PNF_Callback interface, contained in module IA_CL, defines two methods which a PNF server can invoke to either notify the client of "hits" against a standing query or request a location to place an incoming "pushed" image. The PNF_Callback Interface contains two methods as described below:

- **Notify Method -** The notify method is invoked by a PNF server to notify the client that a standing query has new hits. The server will supply the QueryId of the standing query in the first parameter query_identifier and pass the details of the hits in the second parameter results.
- **Push Method** - The push method is invoked by a PNF server when it has imagery or geospatial information to be pushed to the client. The standing query that the imagery or geospatial information resulted from is passed in the first parameter query_identifier and the details of the incoming imagery or geospatial information are passed in the second parameter results. The client is required to return a LocationSpecList containing locations for the incoming imagery or geospatial information in the same order as the results, that is the first LocationSpec in the LocationSPecList corresponds to the first result in the QueryResults etc. For any images that are not wanted, a NULL or empty LocationSpec is returned.

We have judged the approach described above to be the best method currently available to support schema-independent queries in a distributed, client/server-based catalog facility. The OGC should advocate the necessary types and structures in the models and specifications that are developed under its auspices.

3.3 RI3 Interdependencies Among Components

This topic was addressed jointly with "Identified Sets Of Service Components Comprising A Catalog" under RI2.

3.4 RI4 Extensions/Changes To OpenGIS and/or DCP Services Used To Construct Catalogs

We propose no extensions or changes to either the OpenGIS and/or DCP service models.

3.5 RI5 The Proper Role Of Metadata In Cataloging

There is a strong need for standardized metadata in order for catalog facilities to interoperate successfully with one another. NIMA has undertaken a Geospatial Data Standardization Project in order to comply with the Department of Defense directive for its agencies to standardize the data elements used in the computing systems of these agencies as articulated under Directive 8320.1. This Directive also establishes the Defense Data Dictionary System (DDDS) as the central repository for standardized data element and data entities throughout the Department of Defense. The NIMA Geospatial Data Standardization Project serves the DoD's need to populate the DDDS with standard Mapping, Charting, and Geodesy (MC&G) data dictionary information.

The NIMA project initiated a modeling activity to facilitate this standardization process. Two models were produced in IDEF1X form that pertain to metadata and are documented in the "DoD Geospatial Data Standardization Project Report, Volume 3, Geospatial Metadata" [9]. The NIMA models are built on the foundation of predecessor standards such as the FGDC's Content Standard for Digital Geospatial Metadata and the CIO's Standards Profile for Imagery Access (SPIA), and are being actively coordinated with ISO TC 211 as the process of convergence of these standards toward one another continues, as illustrated in Figure 3-3.

The first of these metadata models is the Geospatial Feature Metadata Model that addresses the metadata of a feature, including accuracy, precision, degree of completion, origin, currency, and other metadata needed to manage basic information that describe a feature.

The Geospatial Feature Metadata Model is based on the structures of the Feature and Attribute Coding Catalog (FACC) as defined in Part 4 of the Digital Geographic Information Exchange Standard (DIGEST). DIGEST is an international exchange standards development effort comprised of the 20 nations of the Digital Geographic Information Working Group (DGIWG) that have collaborated to define a coding standard for MC&G features and attributes. The Geospatial Feature Metadata Model exhibits five views relating to attribute data, geometry, management data, topology, and order fulfillment. Of these, only the Geometry view has been approved for inclusion in the DDDS.

The second of these models is the Metadata Data Model for Geospatial Datasets, which is not yet complete. It is based largely on the Federal Geographic Data Committee's (FGDC) Content Standards for Digital Geospatial Metadata. The Metadata Data Model for Geospatial Datasets contains information about data that are needed to control, manage, retrieve, and use or determine the usefulness of geospatial information for a specific requirement. In addition to the seven major areas defined in the FGDC Content Standard For Digital Geospatial

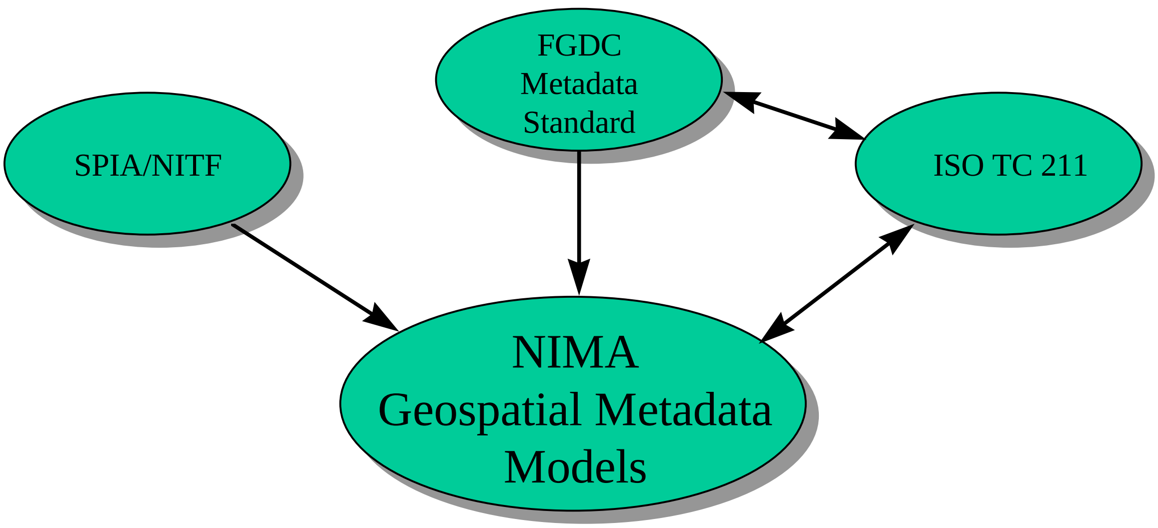

Figure 3-3: Standard Metadata Convergence [2]

Metadata, the Metadata Data Model for Geospatial Datasets extends into domain areas where the FGDC Metadata Standard is deficient for NIMA's requirements, such as imagery and security, and is very similar in that respect to the SPIA. The metadata views that are part of this model are:

- Identification & Citation
- Point Of Contact
- Data Quality
- Spatial Data Representation
- Spatial Reference
- Feature And Attribute
- Distribution
- Metadata Reference
- Image
- Security

The model is currently being synchronized with the metadata work that is associated with ISO TC 211 and will become a profile of the international standard. Additional work to be completed on this model include mechanisms for compound datasets and series of datasets, ensured compliance with DoD Directive 8320.1, mapping to legacy systems and to other standards, and harmonizing this model with the Geospatial Feature Metadata Model. Completion of these items is expected by January 1997.

As described in the IAS Specification [8], guidance on metadata requirements is provided for the USIGS based on the philosophy that metadata comprise all the attribute information in the system. This includes the attributes from the Standards Profile for Imagery Access (SPIA), the Exploitation Support Data (ESD), the NITF headers, the NITF subheaders, and the NITF extension data. The required catalog attributes are identified in the Common Imagery Interoperability Profile (CIIP). These metadata are required to:

- Support a robust catalog search capability.
- Support array (tile) retrieval needs for header file information.
- Isolate stable interface specifications from changing metadata attributes and requirements.
- Eliminate hidden or inconsistent interfaces that access specialty metadata and/or header data.
- Achieve simplicity and uniformity (by making the access to all metadata consistent, it simplifies the interfaces and their usage, thus saving development costs and making interoperability possible.
- Enable access and search of product data embedded in a compound file such as NITF.

The metadata will also contain descriptive information about itself, including an on-line representation of its own schema to allow clients to discover the current attributes supported. To support client update of the library, it is also necessary to store metadata about available distributed libraries which can be referred to through opaque object references. The metadata also contain attributes in support of array retrieval, including the number of levels of resolution, reduced resolution image sizes, and so forth. The inclusion of this metadata does not imply that all metadata is equally capable of being queried. However, it is intended that the access APIs and the schema information be consistent for all forms of system metadata. This is a key design requirement for interoperability.

Besides the catalog, there is an additional set of metadata capabilities provided by the Parameters Interface. The Parameters metadata may include additional information directly associated with specific objects, or state information related to specific clients. The Parameters Interface is an abstract interface inherited by other interfaces. It has the capability to retrieve and change the parameter values that are closely related to particular objects and particular client interactions with these objects. For example, the parameters can include metadata that describes the object. Parameters can also include some controllable client-specific characteristics that are used to modify the behavior of other operations, such as retrieval of array regions.

We recommend a different view than the arbitrary taxonomy of metadata that is provided in the subsections of Request Item 5 in the RFI. Metadata organization should be based on a model developed with sound and proven modeling techniques. Emphasis should be placed on the interfaces necessary to support the model and not on an arbitrary taxonomy such as the one presented in the RFI.

3.6 RI6 Catalog Signature

As described in the IAS Specification [8], there is a locator service defined in the CIIP that supports a multi-faceted capability for managing client access to multiple image libraries. Part of the capabilities of the locator service are transparent to the client; relating to the automatic routing of retrieval requests to alternate libraries. These capabilities are part of the library implementation and are not within the scope of an interface specification.

Another locator service capability addresses client selection of libraries. In this case, there is system metadata information exposed to the client, and this would expose additional client interfaces. There is a commercial standard that can provide this service called a Trader Service.

The Trader Service originated with the Open Distributed Processing standards work at ISO. This work resulted in the fast-track adoption of ISO IDL as the standard method to define software interface bindings for formal standards. The work continued at the OMG to create a commercial API for the Trader Service. This work is still in progress and is expected to be completed in early 1997.

The Trader Service is a yellow-pages directory service. Service offerers can advertise their capabilities via the Trader Services to enable discovery by clients. In their advertisements, service offerers include their ISO IDL interface type and various characteristics that define and discriminate their services.

The Trader Service is highly applicable to the needs of the imagery and geospatial information Locator Service. Using the Trader Service, image libraries can advertise their service location, imagery or geospatial information coverages, and other characteristics. Clients can select the appropriate imagery or geospatial information library (IPL, CIL, NIL, etc.) based upon these characteristics. Each library can be uniquely referenced using an object reference obtained from the Trader Service. The client can select the appropriate library based upon the choice of library object reference when using the Image Access Services APIs.

3.7 RI7 Access Services And Issues

While we recognize that the body of issues presented under Request Item 7 as a group are important issues to information processing, most address lower-level commodity and infrastructure issues that OGC should not be concerned with. For instance, value and referential integrity mechanisms are now built into the COTS database engine products available from the major DBMS vendors. This is an implementation detail that is not germane to the problems that OGC is trying to solve. We recommend that OGC not devote time and resources to such implementation details and focus on services that are inherently and uniquely geospatial in nature. Specific infrastructure-related disconnects that may be identified during the OGIS development process should be forwarded as inputs to OMG and the Open Group as appropriate. Our positions on selected Request Item topics are provided below.

The IAS Specification [8] addresses reliability and error handling issues at a generic interface level. The CAF defines ISO IDL that specifically establishes an ExceptionInfo symbol. This symbol is used in all of the exception definitions in a facility to give these exceptions a consistent set of return values. User exceptions are defined for all cases of bad input parameters as well as error conditions which are unique to particular operations. There is also a set of standard exceptions defined that cover most generic error conditions such as communication failure and lack of access permission. The standard CORBA exceptions specified by OMG in standard ISO IDL are used globally for all facilities of the CIIF. Specifically, there are five additional exceptions that are returned by the CAF operations as shown in the response to Request Item 2.

A BadQuerySyntax exception indicates that the query was improperly formed according to the Boolean Query Syntax rules described under Request Item 8 below. A BadAttribute exception indicates that one or more of the attributes used in the query are inappropriate or unknown. A BadQueryValue exception indicates that a literal-constant value used in the query expression was inappropriate. A BadEllipse exception indicates that the axes or azimuth are inappropriate. A BadQueryId exception is returned if a query identifier is invalid. In addition to these semantics, the ExceptionInfo may return supplementary information which provides further indications of the problem causing the exception.

The CIIF Reference Model [2], addresses the need for integrated security services that are tailored to a distributed computing architecture. Currently there are no existing security-related government regulations or guidance for object-oriented systems. Though the CIIF Reference Model identifies these security services, the specifics of some of them have not yet been adequately addressed.

3.8 RI8 Query Requirements On Catalogs

The approach used to specify the CAF in the IAS Specification [8] emphasizes the definition of interfaces through the use of standard APIs. Standard APIs derive their attractiveness for use from their technology-neutral and schema-independent nature. When APIs are used, implementation issues are not of concern to the system architecture on which the catalog facility application is based. Implementation mechanisms are constantly changing. APIs isolate the application from these constantly changing implementation mechanisms, a characteristic which is especially important in a heterogeneous environment.

3.9 RI9 Catalog Types

A realistic view of imagery and geospatial information catalog implementations recognizes that the world is comprised of a multitude of heterogeneous libraries of digital holdings, each physically distributed and accessible via wide-area network mechanisms. These libraries are owned and maintained by different communities and organizations, and are based on differing physical and logical schemas. A catalog facility such as the CAF would operate equally well on local catalog holdings as it would on global catalog holdings, regardless of whether these catalog holdings are centralized or distributed. As described in Request Item 8, implementation of the support mechanisms for catalog services in the CAF is an infrastructure problem for the server and is not an issue for the client. A catalog server of this type may be expected to be somewhat challenging to implement.

Individual catalogs themselves would be registered with trader services to enable their discovery. OMG is expected to release an RFP on Internet Information Discovery based on IIOP in the January-March 1997 timeframe that may shed some new light on these issues. It is recommended that OGC consult OMG for new or original ideas on uniform or mixed catalog content, and the implications of these for catalog discovery.

3.10 RI10 Naming Issues

Naming and persistence issues are handled in the USIGS as part of the Object Services component of the USIGS, as described in the CIIF Reference Model [2]. Object Services standardize the life-cycle management of objects. Interfaces are provided to create objects, to control access to objects, to keep track of relocated objects, and to control the relationship between styles of objects through class management mechanisms. Also provided are the generic environments in which single objects can perform their tasks. Object services provide for application consistency and help to increase programmer productivity. Please refer to the CIIF Reference Model diagram in Section 2 to determine their place in relation to the CAF and other components.

A naming service is a key component of the USIGS Technical Architecture. This service provides the ability to bind a name to an object relative to a naming context. A naming context is an object that contains a set of name bindings in which each name is unique. To resolve a name is to determine the objects associated with the name in a given context. Through the use of a “names library,” name manipulation is simplified and names can be made representation independent, thus allowing their representation to evolve without requiring client changes.

A persistent object service is also a key component of the USIGS Technical Architecture. This service provides common interfaces to the mechanisms used for retaining and managing the persistent state of objects. This service is to be used in conjunction with other object services, for example, naming, relationships, transactions, life cycle, etc. This service has the primary responsibility for storing the persistent state of objects, with other services providing other capabilities.

The Object Services approach articulated by the CIIF Reference Model is not unique to the USIGS and reflects current developments on naming and persistence issues that are occurring in the OMG. For issues such as these that are primarily related to infrastructure and are not unique to imagery and geospatial information, OGC should rely on the standard approaches emerging from OMG and other standards bodies and consortia for the development of the necessary interoperability mechanisms. OGC should examine the output from these organizations to determine any aspects of these mechanisms that are deficient for imagery and geospatial information, and then collaborate with those responsible parties to remedy any deficiencies that may be found for these types of information.

3.11 RI13 Information Concerning Existing Implementations

The questions in the request item address very specific lower-level infrastructure details for catalog service implementations. The services and interfaces of the CAF are defined at a higher level to be technology-neutral and schema-independent, as stated in RI8. The CAF is capable of being implemented using a variety of standard implementation approaches. It is very likely that some implementations of the CAF may use Z39.50, especially since the use of the Government Information Locator Service (GILS) has been mandated for use by the Federal Government for catalog support services. GILS is based on a Z39.50 profile. Other implementations of the CAF may possibly use HTTP, IIOP, ESIOP, DCE/RPC, DCOM, RMI, or some other set of protocols that may evolve in the future.

We agree that the questions asked in this request item are valid and of importance to catalog service implementers from an infrastructure and implementation perspective. OGC, however, should not concern itself with the specific details of these implementations. Instead, OGC should keep abreast of developments in the named areas to determine any aspects of these implementation mechanisms that are deficient for imagery and geospatial information, and then collaborate with those responsible parties to remedy any deficiencies that may be found for these types of information, as previously stated. OGC, in evaluating existing catalog service implementations, should focus on identifying common services, interfaces, and architectural components, and ensuring that they are consistent with the OGIS Abstract Specification.

3.12 RI14 Distributed Computing Platform (DCP) Options

The implementations listed in the request item are language independent and have no impact on the interfaces specified for a catalog service facility. As long as the interfaces are written in ISO IDL, the details of implementation do not matter.

ISO IDL was chosen as the specification mechanism for the CIIF because it offers many inherent advantages over older API specification methods. ISO IDL is a formal language (similar in appearance to a C++ header file) that is used to define the interfaces between interoperable software objects. Among its chief advantages is that ISO IDL can be directly compiled into any of several common programming languages, using standard language mappings, to automatically set up the mechanisms needed to pass service requests across the network in a distributed software architecture. By adopting ISO IDL to develop API specifications, the CIIF will capitalize on the many interoperability benefits inherent to the use of ISO IDL. These include:

- ISO IDL is a widely accepted specification language for modern ICDs. It is actively supported by major standards bodies such as the International Organization of Standards (ISO, Draft International Standard 14750 [3]), X/Open, OMG, and OGC. ISO IDL has emerged as an accepted standard for specifying other standards.
- ISO IDL is vendor, platform, and language independent. A single specification suffices for C, C++, Ada, and Smalltalk, which reduces the cost and complexity of documentation while promoting improved specification rigor.
- ISO IDL is a complete and rigorous specification notation. It features strong typing, standard language mappings, exception handling, etc.
- ISO IDL's standard language bindings enable an ISO IDL specification to be automatically compiled into any of the standard programming languages. The result can be used both to provide an initial framework for software development, and to provide the basis for comparing as-built code to original specifications.
- ISO IDL can be used with, or independently of, OMG's Object Management Architecture.

As previously recommended, OGC should focus on identifying common services, interfaces, and architectural components, and on ensuring that the OGIS Abstract Specification accommodates these.

4. List Of References

List of Documents Referenced in the Text of the Request Item Responses:

[1] Global Geospatial Information and Services (GGIS) for the Warrior, National Imagery and Mapping Agency, http://164.214.2.59/cgi-bin/waisgate?WAISdocID=087986451+0+0+0&WAISaction=retrieve, 4 March 1995.

[2] Common Imagery Interoperability Facilities (CIIF) Reference Model, Version 2.0; National Imagery and Mapping Agency (CIO-2061); 20 December 1996.

[3] Information Technology - Open Distributed Processing - Interface Definition Language, Draft Edition 1; International Organization of Standards (ISO/IEC DIS 14750), date not available.

[4] Reference Model For An Open Archival Information System (OAIS), Draft Version 7.0; International Organization of Standards (MUN/96/P2/N11), http://www.gsfc.nasa.gov/nost/isoas/overview.html; 25 October 1996.

[5] IEEE Reference Model for Open Storage Systems Interconnection (OSSI); www.arl.mil/IEEE/ssswg.html, multiple volumes, various dates.

[6] "Z39.50 Profile for Access to Digital Collections"; CNIDR, http://vinca.cnidr.org/protocols/profiles/zdl.html

[7] Z39.50 Profile for Access to Digital Library Objects; Library of Congress, http://lcweb.loc.gov/z3950/agency/profiles/dl.html, August 15, 1996.

[8] Imagery Access Services (IAS) Specification, United States Imagery System, Version 1.1; National Imagery and Mapping Agency (CIO-2068); 20 December 1996.

[9] DoD Geospatial Data Standardization Project Report, Volume 3, Geospatial Metadata; National Imagery and Mapping Agency, 16 September 1996.

List of Other Documents of Relevance:

Joint Technical Architecture (JTA), Version 1.0; Department of Defense (DoD), http://www.itsi.disa.mil/jta/, 22 August 1996.

The Common Object Request Broker: Architecture and Specification (CORBA) Version 2.0; Object Management Group (OMG) (PTC/96-03-04), July 1995.

The Portable Operating System Interface (POSIX); Institute of Electrical and Electronics Engineers, Inc. (IEEE P1003.0), multiple volumes, various dates.

DoD Interoperable Map Software (IMS) MC&G Domain Definition Report, Version 2; Department of Defense, 1 June 1996.